Update or Wait: How to Keep Your Data Fresh

Yin Sun[†], Elif Uysal-Biyikoglu[‡], Roy D. Yates[⋆], C. Emre Koksal[†], and Ness B. Shroff[†§]
[†]Dept. of ECE, [§]Dept. of CSE, The Ohio State University, Columbus, OH
[‡]Dept. of EEE, Middle East Technical University, Ankara, Turkey
[⋆]Dept. of ECE, Rutgers University, North Brunswick, NJ

arXiv:1601.02284v8 [cs.IT] 26 Sep 2016

***Abstract*—In this work we study how to manage the freshness of status updates sent from a source to a destination via a communication server. A proper metric of data freshness at the destination is the age-of-information, which is defined as how old the freshest received update is since the moment this update was generated at the source. A logical policy is the *zero-wait* policy, i.e., the source submits a fresh update once the server is free, which achieves the maximum throughput and the minimum average delay. Surprisingly, this zero-wait policy does not always minimize the average age. This motivates us to study how to optimally control the status updates to keep data fresh and to understand when the zero-wait policy is optimal. We introduce a penalty function to characterize the level of dissatisfaction on data staleness and formulate the average age penalty minimization problem as a constrained semi-Markov decision process (SMDP) with an uncountable state space. Despite the difficulty of this problem, we develop efficient algorithms to find the optimal status update policy and provide sufficient and necessary conditions to characterize the optimality of the zero-wait policy. We show that the zero-wait policy can be far from the optimum if (i) the penalty function grows quickly with respect to the age, or (ii) the update service times are positively correlated and follow a heavy-tail distribution. To the best of our knowledge, this is the first optimal control policy which is proven to minimize the age-of-information in status update systems.**

I. Introduction

In recent years, the proliferation of mobile devices and applications has significantly boosted the need for real-time information updates, such as news, weather reports, email notifications, stock quotes, social updates, mobile ads, etc. Timely status updates are also critical in network-based monitoring and control systems, including sensor networks used in temperature and air pollution monitoring [1], surround monitoring in autonomous vehicles [2], and phasor data updates in power grid stabilization systems [3].

A common need in these applications is to maximize the freshness of the data at the destination of information updates. In light of this, a metric called the *age-of-information*, or simply *age*, was defined in [4]. At time t, if the freshest update at the destination has a time stamp $U(t)$, the age is

$$\Delta(t) = t - U(t). \tag{1}$$

Hence, the age is the time elapsed since the freshest received packet was generated.

This paper was presented in part at IEEE INFOCOM 2016.

This work was supported in part by DTRA grant HDTRA1-14-1-0058, NSF grants CNS-1446582, CNS-1409336, CNS-1518829, CNS-1514260, CNS-1422988, and CNS-1054738, ARO grant W911NF-14-1-0368, and ONR grant N00014-15-1-2166. E. Uysal-Biyikoglu was supported in part by TUBITAK and in part by a Science Academy BAGEP award.

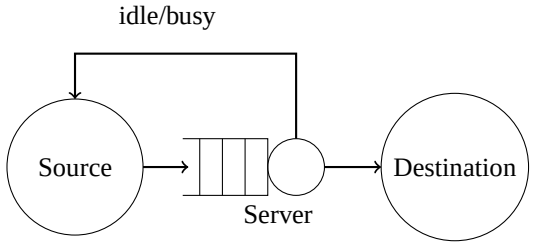

Fig. 1. System model.

Most existing research on the age-of-information focuses on an "enqueue-and-forward" model [4]–[10], where status updates are randomly generated or arrive at a source node. The source enqueues these updates and forwards them later to a remote destination through a network. It is worth noting that the goal of age minimization differs from those of throughput maximization and delay minimization: A high update frequency increases the system throughput, but may also induce a large waiting time in the queue which in turn increases the age; on the other hand, a low update frequency can reduce the queueing delay, but the destination may end up having stale status information due to not enough updates [4], [8], [10]. In [4], [9], it was found that a good policy is to discard the old packets waiting in the queue if a new sample arrives, which can greatly reduce the impact of queueing delay.

In this paper, we study a "generate-at-will" model depicted in Fig. 1. In this model, the source monitors the server's idle/busy state, and in contrast to [4]–[10], is able to generate status updates at any time by its own will. Hence, no updates need to be generated when the server is busy, which completely eliminates the waiting time in the queue and hence the queue in Fig. 1 is always empty. A simple *zero-wait* policy, also known as the *work-conserving* policy in queueing theory, that submits a fresh update once the server becomes idle, achieves the maximum throughput and the minimum average delay. Surprisingly, this zero-wait policy does **not** always minimize the average age of the information [11]. The following example reveals the reason behind this phenomenon:

Example: *Suppose that the source submits a stream of update packets to a remote destination. The service times of these updates form a periodic sequence*

$$0, 0, 2, 2, 0, 0, 2, 2, 0, 0, 2, 2, \ldots$$

Suppose that update 1 *is generated and submitted at time* 0 *and delivered at time* 0*. Under the zero-wait policy, update* 2 *is also generated at time* 0 *and delivered at time* 0*. However, despite its negligible service time, update* 2 *has not brought any new*

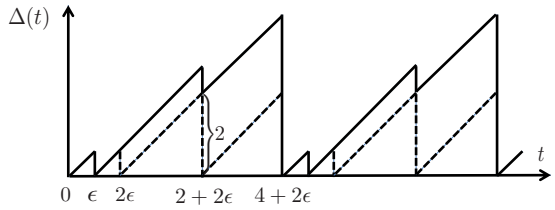

Fig. 2. Evolution of the age $\Delta(t)$ in the ϵ-wait policy in the example.

information to the destination after update 1 *was delivered, because both updates were sampled at time* 0*. Therefore, the potential benefit of the zero service time of update* 2 *is wasted! This issue occurs periodically over time: Whenever two consecutive updates have zero service time, the second update of the two is wasted. Therefore, a 1/4 of the updates in this sequence are wasted in the zero-wait policy!*

For comparison, consider the ϵ*-wait policy that waits for* ϵ *seconds after each update with a zero service time, and does not wait after each update with a service time of* 2 *seconds. The time-evolution of the age* $\Delta(t)$ *in the non-zero-wait policy is shown in Fig. 2. Update* 1 *is generated and delivered at time* 0*. Update* 2 *is generated and delivered at time* ϵ*. Update* 3 *is generated at time* 2ϵ *and is delivered at time* $2+2\epsilon$*. Because the service time of update* 3 *is* 2 *seconds, the latest delivered update at time* $2+2\epsilon$ *is of the age* 2 *seconds. Hence, the age* $\Delta(t)$ *drops to* 2 *seconds at time* $2+2\epsilon$*. Update* 4 *is generated at time* $2+2\epsilon$ *and is delivered at time* $4+2\epsilon$*. At time* $4+2\epsilon$*, the age drops to zero because update* 5 *is generated at this time and is delivered immediately.*

It can be shown that the time-average age of the ϵ*-wait policy is*

$$\begin{aligned}&(\epsilon^2/2+\epsilon^2/2+2\epsilon+4^2/2)/(4+2\epsilon)\\=&(\epsilon^2+2\epsilon+8)/(4+2\epsilon)\ \textit{seconds.}\end{aligned}$$

If the waiting time is $\epsilon=0.5$*, the time-average age of the* ϵ*-wait policy is* 1.85 *seconds. If the waiting time is* $\epsilon=0$*, it reduces to the zero-wait policy, whose time-average age is* 2 *seconds. Hence, the zero-wait policy is not optimal!*

This example point out a key difference between data communication systems and status update systems: In communication systems, all packets are equally important; however, in status update systems, an update packet is useful only if it carries some fresh information to the destination. While the theory of data communications is quite mature, the optimal control of status updates remains open.

For a source that can insert waiting times between updates, the aim of this paper is to answer the following questions: How to optimally submit update packets to maximize data freshness at the destination? When is the zero-wait policy optimal? To that end, the following are the key contributions of this paper:

- We generalize existing status update studies by introducing two new features: age penalty functions and correlated (non-*i.i.d.*) service processes. We define an age penalty function $g(\Delta)$ to characterize the application-specific level of dissatisfaction for data staleness, where $g(\cdot)$ is *measurable, non-negative, and non-decreasing*. The update service process is modeled as a stationary ergodic Markov chain with an uncountable state space, which generalizes the *i.i.d.* service processes studied in previous work [4]–[12].
- We formulate the average age penalty minimization problem as a constrained semi-Markov decision process (SMDP) with an uncountable state space. Despite the difficulty of this problem, we manage to solve it by a divide-and-conquer approach: We first prove that there exists a stationary randomized policy that is optimal for this problem (Theorem 1). Further, we prove that there exists a stationary deterministic policy that is optimal for this problem (Theorem 2). Finally, we develop a low-complexity algorithm to find the optimal stationary deterministic policy (Theorem 3). To the best of our knowledge, this is the first optimal control policy which is proven to minimize the age-of-information (i.e., maximize data freshness) in status update systems.
- We further investigate when the zero-wait policy is optimal. For the special case of proportional penalty function and *i.i.d.* service times, we devise a simpler solution to minimize the average age (Theorem 4), and obtain a sufficient and necessary condition to characterize when the zero-wait policy is optimal (Theorem 5). For general age penalty functions and correlated service processes, several sufficient conditions for the optimality of the zero-wait policy are provided (Lemma 3).
- Our theoretical and simulation results demonstrate that, in many scenarios, the optimal policy is to wait for a certain amount of time before submitting a new update. In particular, the zero-wait policy can be far from optimality if (i) the penalty function grows quickly with respect to the age, or (ii) the update service times are positively correlated and follow a heavy-tail distribution.

II. Model and Formulation

A. System Model

We consider a status update system depicted in Fig. 1, where a source generates update packets and sends them to a remote destination through a communication server. This server provides First-Come First-Served (FCFS) service to the submitted update packets. The service of an update packet is considered complete, when it is successfully received by the destination. After that, the server becomes available for sending the next packet.

The source generates and submits updates at successive times $S_0, S_1, \ldots$ Update i, submitted at time S_i, is delivered at time $D_i = S_i + Y_i$, where $Y_i \geq 0$ is the service time of update i. Suppose that update 0 is submitted to an idle server at time $S_0 = -Y_0$ and delivered at $D_0 = 0$, as shown in Fig. 3. The source has access to the idle/busy state of the server and is able to generate updates at any time by its own will. Hence, the source should not generate update $i+1$ when the server is busy processing update i, because this will incur an unnecessary waiting time in the queue. After update i is delivered at time D_i, the source may introduce a *waiting time* $Z_i \in [0, M]$ before submitting update $i+1$ at time $S_{i+1} = D_i + Z_i$, where M is defined as a maximum

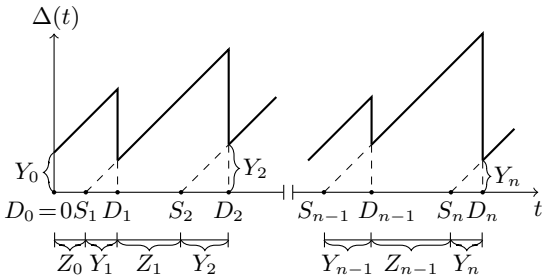

Fig. 3. Evolution of the age-of-information $\Delta(t)$.

waiting time allowed by the system. The source can switch to a low-power sleep mode during $[D_i, S_{i+1})$. We assume that the service process $(Y_0, Y_1, \ldots)$ is *a stationary and ergodic Markov chain* with a possibly uncountable state space and a positive mean $0 < \mathbb{E}[Y_i] < \infty$, which generalizes the *i.i.d.* service processes in previous studies [4]–[12].[1] This Markovian service process model is introduced to study the impact of temporal-correlation on the optimality of the zero-wait policy. In Section IV, we will see that the zero-wait policy is close to the optimum when the service process is negatively correlated; but can be far from the optimum when the service process is positively correlated.

At any time t, the most recently received update packet is time-stamped with

$$U(t) = \max\{S_i : D_i \leq t\}. \tag{2}$$

The *age-of-information* $\Delta(t)$ is defined as [4]

$$\Delta(t) = t - U(t), \tag{3}$$

which is also referred to as *age*. As shown in Fig. 3, the age $\Delta(t)$ is a stochastic process that increases linearly with t between updates, with downward jumps occurring when updates are delivered. Specifically, when update i is sent at time $t = S_i$, it is delivered at time $D_i = S_i + Y_i$ with age $\Delta(D_i) = D_i - S_i = Y_i$. After that, the age increases linearly and reaches $\Delta(D_{i+1}^-) = Y_i + Z_i + Y_{i+1}$ just before update $i+1$ is delivered. Then, at time D_{i+1}, the age drops to $\Delta(D_{i+1}) = Y_{i+1}$.

B. Problem Formulation

We introduce an age penalty function $g(\Delta)$ to represent the level of dissatisfaction for data staleness or the need for new information update. The function $g : [0, \infty) \to [0, \infty)$ is assumed to be *measurable, non-negative, and non-decreasing*. Some examples of $g(\cdot)$ are the power function $g(\Delta) = \Delta^a$, the exponential function $g(\Delta) = e^{a\Delta}$, and the stair-step function $g(\Delta) = \lfloor a\Delta \rfloor$, where $a \geq 0$ and $\lfloor x \rfloor$ is the largest integer no greater than x. Two age penalty functions are depicted in Figure 4. Note that this age penalty model is quite general, as it allows $g(\cdot)$ to be discontinuous and non-convex.

To analyze the average age penalty, we decompose the area under the curve $g(\Delta(t))$ into a sum of disjoint components:

[1]The results in this paper can be readily extended to a more general Markovian service process $(Y_0, Y_1, \ldots)$ with a longer memory, in which W_i is defined as $W_i = (Y_i, Y_{i+1}, \ldots, Y_{i+k})$ for some finite k, and the sequence $(W_0, W_1, \ldots)$ forms a Markov chain.

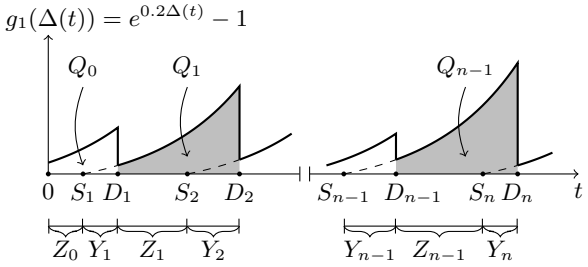

(a) An exponential age penalty function $g_1(\Delta) = e^{0.2\Delta} - 1$.

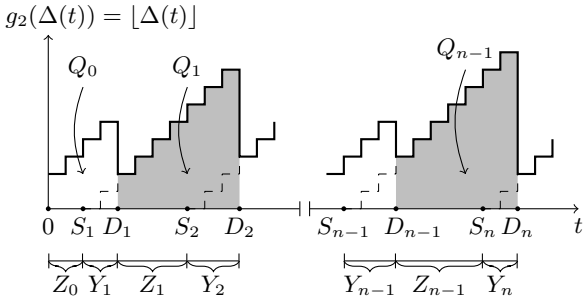

(b) A stair-step age penalty function $g_2(\Delta) = \lfloor \Delta \rfloor$.

Fig. 4. Evolutions of two age penalty functions.

Consider the time interval $[0, D_n]$, where $D_n = \sum_{i=0}^{n-1}(Z_i + Y_{i+1})$. In this interval, the area under $g(\Delta(t))$ can be seen as the concatenation of the areas Q_i, $0 \leq i \leq n-1$, such that

$$\int_0^{D_n} g(\Delta(t))dt = \sum_{i=0}^{n-1} Q_i,$$

where

$$Q_i = \int_{D_i}^{D_i+Z_i+Y_{i+1}} g(\Delta(t))dt = \int_{Y_i}^{Y_i+Z_i+Y_{i+1}} g(\tau)d\tau. \tag{4}$$

In the second equation of (4), we have used the fact that $\Delta(t) = t - S_i = t - (D_i - Y_i)$ for $t \in [D_i, D_i + Z_i + Y_{i+1}]$. Let us define

$$q(y, z, y') = \int_y^{y+z+y'} g(\tau)d\tau. \tag{5}$$

Then, Q_i can be expressed as $Q_i = q(Y_i, Z_i, Y_{i+1})$. Since $g(\Delta)$ is non-negative, the function $q(y, z, y')$ is non-decreasing in z. We assume

$$\mathbb{E}\left[q(Y_i, M, Y_{i+1})\right] < \infty, \tag{6}$$

implying $\mathbb{E}\left[q(Y_i, Z_i, Y_{i+1})\right] < \infty$ for all $Z_i \in [0, M]$.

Our goal is to minimize the average age penalty by controlling the sequence of waiting times $(Z_0, Z_1, \ldots)$. Let $\pi \triangleq (Z_0, Z_1, \ldots)$ denote a status update policy. We consider the class of *causally feasible* policies, in which control decisions are made based on history and current information of the system, as well as the distribution of the service process $(Y_0, Y_1, \ldots)$. Specifically, Z_i is determined based on the past realizations of $(Y_0, Y_1, \ldots, Y_i)$, without using the realizations of future service times $(Y_{i+1}, Y_{i+2}, \ldots)$; but the conditional distribution of $(Y_{i+1}, Y_{i+2}, \ldots)$ given $(Y_0, Y_1, \ldots, Y_i)$ is available. Let Π denote the set of all causally feasible policies satisfying $Z_i \in [0, M]$ for all i.

The average age penalty per unit time is defined by[2]

$$\limsup_{n\to\infty} \frac{\mathbb{E}\left[\int_0^{D_n} g(\Delta(t))dt\right]}{\mathbb{E}\left[D_n\right]}. \tag{7}$$

Because the service process $(Y_0, Y_1, \ldots)$ is stationary and ergodic, we can obtain $\mathbb{E}[Y_i] = \mathbb{E}[Y_{i+1}]$ and hence $\mathbb{E}\left[D_n\right] = \mathbb{E}[\sum_{i=0}^{n-1}(Z_i + Y_{i+1})] = \mathbb{E}[\sum_{i=0}^{n-1}(Y_i + Z_i)]$. Using this, the dynamic control problem for minimizing the average age penalty can be formulated as

$$\overline{g}_{\text{opt}} = \min_{\pi\in\Pi} \limsup_{n\to\infty} \frac{\mathbb{E}\left[\sum_{i=0}^{n-1} q(Y_i, Z_i, Y_{i+1})\right]}{\mathbb{E}[\sum_{i=0}^{n-1}(Y_i + Z_i)]} \tag{8}$$

$$\text{s.t.} \quad \liminf_{n\to\infty} \frac{1}{n}\mathbb{E}\left[\sum_{i=0}^{n-1}(Y_i + Z_i)\right] \geq \frac{1}{f_{\max}}, \tag{9}$$

where $\overline{g}_{\text{opt}}$ is the optimum objective value of Problem (8), the expectation $\mathbb{E}$ is taken over the stochastic service process $(Y_0, Y_1, \ldots)$ for given policy π, and $f_{\max}$ is the maximum allowed average sampling rate due to a long-term average resource constraint (i.e., on the power resource or CPU cycles) spent on generating status updates. We assume $M > 1/f_{\max}$ such that Problem (8) is always feasible and $\overline{g}_{\text{opt}} < \infty$. In this paper, we will study Problem (8) both with and without the constraint (9). In Section IV, sufficient and necessary conditions will be provided to characterize when the zero-wait policy is optimal for solving Problem (8) without the constraint (9).

In Problem (8), Y_i is the state of an embedded Markov chain, Z_i is the control action taken after observing Y_i, $Y_i + Z_i$ is the period of stage i, and $q(Y_i, Z_i, Y_{i+1})$ is the reward related to both stage i and $i+1$. Therefore, Problem (8) belongs to the class of constrained semi-Markov decision processes (SMDP) with a possibly uncountable state space, which is generally known to be quite difficult. The class of SMDPs include Markov decision problems (MDPs) [17], [18] and optimization problems of renewal processes [19] as special cases. Most existing studies on SMDPs deal with (i) unconstrained SMDPs, e.g., [13], [17], [20], [21], or (ii) constrained SMDPs with a countable state space, e.g., [16], [22]–[24]. However, the optimality equations (e.g., Bellman's equation) for solving unconstrained SMDPs are not applied to constrained SMDPs [25], and the studies on constrained SMDPs with a countable state space cannot be directly applied to Problem (8) which has an uncountable state space.

[2]There are two widely used definitions of average cost per unit time in infinite-horizon undiscounted semi-Markov decision processes (SMDP) [13]–[17]: In one definition, the average cost is the limit of expected total cost over a finite deterministic horizon divided by the length of the horizon, i.e., $\limsup_{T\to\infty} \frac{1}{T}\mathbb{E}\left[\int_0^T g(\Delta(t))dt\right]$; in the second definition, the average cost is the limit of the expected total cost over a finite number of stages divided by the expected cumulative time of these stages, as in (7). These two definitions are both reasonable [13], [16], [17]. They are equal under stationary randomized policies in which the generated semi-Markov chain has one ergodic class; see [13]–[15] for finite and countable state models. In general, however, these criteria are different. In our study, the second definition turns out to be analytically convenient.

III. Optimal Status Update Policy

In this section, we develop a chain of theoretical results to solve Problem (8) by exploiting its structure: First, we prove that there exists a *stationary randomized* policy that is optimal for Problem (8). Further, we prove that there exists a *stationary deterministic* policy that is optimal for Problem (8). Finally, we develop a low-complexity algorithm to find the *optimal* stationary deterministic policy that solves Problem (8).

A. Optimality of Stationary Randomized Policies

A policy $\pi \in \Pi$ is said to be a *stationary randomized* policy, if it observes Y_i and then chooses a waiting time $Z_i \in [0, M]$ based only on the observed value of Y_i. In this case, Z_i is determined according to a conditional probability measure $p(y, A) \triangleq \Pr[Z_i \in A | Y_i = y]$ that is invariant for all $i = 0, 1, \ldots$ We use Π_{SR} ($\Pi_{\text{SR}} \subseteq \Pi$) to denote the set of stationary randomized policies such that

$$\Pi_{\text{SR}} = \{\pi \in \Pi : \text{Given the observation } Y_i = y_i,\ Z_i \text{ is chosen according to the probability measure } p(y_i, A) \text{ for all } i\}.$$

Note that (Y_i, Z_i, Y_{i+1}) is stationary and ergodic for all stationary randomized policies. In the sequel, when we refer to the stationary distribution of a stationary randomized policy $\pi \in \Pi_{\text{SR}}$, we will remove subscript i. In particular, the random variables (Y_i, Z_i, Y_{i+1}) are replaced by (Y, Z, Y'), where Z is chosen based on the conditional probability measure $\Pr[Z \in A | Y = y] = p(y, A)$ after observing $Y = y$, and (Y, Y') have the same joint distribution as (Y_0, Y_1). The first key result of this paper is stated as follows:

Theorem 1: (Optimality of Stationary Randomized Policies) If $g : [0, \infty) \to [0, \infty)$ is measurable and non-negative, $(Y_0, Y_1, \ldots)$ is a stationary ergodic Markov chain with $Y_i \geq 0$ and $0 < \mathbb{E}[Y_i] < \infty$, and condition (6) is satisfied, then there exists a stationary randomized policy that is optimal for Problem (8).

Proof sketch of Theorem 1: For any policy $\pi \in \Pi$, define the finite time-horizon average occupations

$$a_{n,\pi} \triangleq \frac{1}{n}\mathbb{E}\left[\sum_{i=0}^{n-1} q(Y_i, Z_i, Y_{i+1})\right] - \frac{\overline{g}_{\text{opt}}}{n}\mathbb{E}\left[\sum_{i=0}^{n-1}(Y_i + Z_i)\right], \tag{10}$$

$$b_{n,\pi} \triangleq \frac{1}{n}\mathbb{E}\left[\sum_{i=0}^{n-1}(Y_i + Z_i)\right]. \tag{11}$$

Let Γ_{SR} be the set of limit points of sequences $((a_{n,\pi}, b_{n,\pi}), n = 1, 2, \ldots)$ associated with all stationary randomized policies in Π_{SR}. We first prove that Γ_{SR} is convex and compact. Then, we show that there exists an optimal policy π_{opt} of Problem (8), such that the sequence $((a_{n,\pi_{\text{opt}}}, b_{n,\pi_{\text{opt}}}), n = 1, 2, \ldots)$ associated with policy π_{opt} has a limit point (a^*, b^*) satisfying $(a^*, b^*) \in \Gamma_{\text{SR}}$, $a^* \leq 0$, and $b^* \geq \frac{1}{f_{\max}}$. Since $(a^*, b^*) \in \Gamma_{\text{SR}}$, there exists a stationary randomized policy π^* achieving this limit point (a^*, b^*). Finally, we show that policy π^* is optimal for Problem (8), which completes the proof. The details are available in Appendix A. ■

The convexity and compactness properties of the set of average occupation measures are essential in the study of constrained MDPs [26, Sec. 1.5], which dates back to Derman's monograph in 1970 [27]. Recently, it was used in

Algorithm 1 Two-layer bisection method for Problem (13)

given $l=0$, sufficiently large $u>\overline{g}_{\text{opt}}$, tolerance ϵ_1.
repeat
  $c:=(l+u)/2$.
  given $\zeta_l=0$, sufficiently large $\zeta_u>0$, tolerance ϵ_2.
  $\zeta:=\zeta_l,\ \nu:=\zeta+c$.
  Compute $z_\nu(\cdot)$ in (17).
  if $\mathbb{E}[Y+z_\nu(Y)]<\frac{1}{f_{\max}}$ **then**
    repeat
      $\zeta:=(\zeta_l+\zeta_u)/2,\ \nu:=\zeta+c$.
      Compute $z_\nu(\cdot)$ in (17).
      if $\mathbb{E}[Y+z_\nu(Y)]\geq\frac{1}{f_{\max}}$, $\zeta_u:=\zeta$; **else**, $\zeta_l:=\zeta$.
    until $\zeta_u-\zeta_l\leq\epsilon_2$.
  end if
  if $f(c)\leq 0$, $u:=c$; **else**, $l:=c$.
until $u-l\leq\epsilon_1$.
return $z(\cdot):=z_\nu(\cdot)$.

stochastic optimization for discrete-time queueing systems and renewal processes, e.g., [19], [28]. The techniques in these studies cannot directly handle constrained SMDPs with an uncountable state space, like Problem (8). One crucial novel idea in our proof is to introduce $\overline{g}_{\text{opt}}$ in the definition of average occupation in (10), which turns out to be essential in later steps for showing the optimality of the stationary randomized policy π^*. In addition, we have also used one property of Problem (8) in the proof: the observation Y_{i+1} depends only on the immediately preceding state Y_i and not on the history state and control $Y_0,\ldots,Y_{i-1}$, $Z_0,\ldots,Z_{i-1}$.

By Theorem 1, we only need to consider the class of stationary randomized policies Π_{SR}. Therefore, as shown in Appendix A, Problem (8) can be simplified to the following functional optimization problem:

$$\min_{p(y,A)}\ \frac{\mathbb{E}[q(Y,Z,Y')]}{\mathbb{E}[Y+Z]} \tag{12}$$
$$\text{s.t.}\ \ \mathbb{E}[Y+Z]\geq\frac{1}{f_{\max}}$$
$$0\leq Z\leq M,$$

where $p(y,A)=\Pr[Z\in A|Y=y]$ is the conditional probability measure of some stationary randomized policy, and (Y,Y') have the same distribution as (Y_0,Y_1).

B. Optimality of Stationary Deterministic Policies

A policy $\pi\in\Pi_{\text{SR}}$ is said to be a *stationary deterministic* policy if $Z_i=z(Y_i)$ for all $i=0,1,\ldots$, where $z:[0,\infty)\to[0,M]$ is a deterministic function. We use Π_{SD} ($\Pi_{\text{SD}}\subseteq\Pi_{\text{SR}}$) to denote the set of stationary deterministic policies such that

$$\Pi_{\text{SD}}=\{\pi\in\Pi_{\text{SR}}: Z_i=z(Y_i)\ \text{for all}\ i\}.$$

Theorem 2: (Optimality of Stationary Deterministic Policies) If $g:[0,\infty)\to[0,\infty)$ is measurable and non-decreasing, then there exists a stationary deterministic policy that is optimal for Problem (12).

Proof sketch of Theorem 2: Since $g(\Delta)$ is non-decreasing, $q(y,z,y')$ is convex in z for any fixed y and y'. Using Jensen's inequality, we can show that for any feasible stationary randomized policy $\pi_1\in\Pi_{\text{SR}}$, there is a feasible stationary deterministic policy that is no worse than policy π_1. The details are provided in Appendix B. ■

Let μ_Y be the probability measure of Y_i, then any bounded measurable function $z:[0,\infty)\to[0,M]$ belongs to the Lebesgue space $L^2(\mu_Y)$ [29, Section 3], because

$$\int_0^\infty |z(y)|^2 d\mu_Y(y)\leq\int_0^\infty M^2 d\mu_Y(y)=M^2<\infty.$$

By Theorems 1 and 2, we only need to consider the class of stationary deterministic policies Π_{SD} and Problem (8) is simplified to the following functional optimization problem:

$$\min_{z(\cdot)\in L^2(\mu_Y)}\ \frac{\mathbb{E}\left[q(Y,z(Y),Y')\right]}{\mathbb{E}[Y+z(Y)]} \tag{13}$$
$$\text{s.t.}\ \ \mathbb{E}[Y+z(Y)]\geq\frac{1}{f_{\max}} \tag{14}$$
$$0\leq z(y)\leq M,\ \forall\ y\geq 0,$$

where $z(\cdot)$ is the function associated with a stationary deterministic policy $\pi\in\Pi_{\text{SD}}$, and (Y,Y') have the same distribution as (Y_0,Y_1). The optimum objective value of Problem (13) is equal to $\overline{g}_{\text{opt}}$.

C. A Low Complexity Solution to Problem (13)

Lemma 1: If $g:[0,\infty)\to[0,\infty)$ is measurable, non-negative, and non-decreasing, then the functional $h:L^2(\mu_Y)\to[0,\infty)$ defined by

$$h(z)=\frac{\mathbb{E}\left[q(Y,z(Y),Y')\right]}{\mathbb{E}[Y+z(Y)]}$$

is quasi-convex.

Proof: See Appendix C. ■

Therefore, Problem (13) is a functional quasi-convex optimization problem. In order to solve Problem (13), we consider the following functional convex optimization problem with a parameter c:

$$f(c)=\min_{z(\cdot)\in L^2(\mu_Y)}\ \mathbb{E}\left[q(Y,z(Y),Y')\right]-c\mathbb{E}[Y+z(Y)] \tag{15}$$
$$\text{s.t.}\ \ \mathbb{E}[Y+z(Y)]\geq\frac{1}{f_{\max}} \tag{16}$$
$$0\leq z(y)\leq M,\ \forall\ y\geq 0.$$

It is easy to show that $\overline{g}_{\text{opt}}\leq c$ if and only if $f(c)\leq 0$ [30]. Therefore, we can solve Problem (13) by a two-layer nested algorithm, such as Algorithm 1. In the inner layer, we use *bisection* to solve Problem (15) for any given parameter c; in the outer layer, we employ *bisection* again to search for a c^* such that $f(c^*)=0$ and thus $\overline{g}_{\text{opt}}=c^*$. Algorithm 1 has low complexity. It requires at most $\lceil\log_2((u-l)/\epsilon_1)\rceil\times\lceil\log_2((\zeta_u-\zeta_l)/\epsilon_2)\rceil$ iterations to terminate and each iteration involves computing $\mathbb{E}[z_\nu(Y)]$ based on (17). The optimality of Algorithm 1 is guaranteed by the following theorem:

Theorem 3: If $g:[0,\infty)\to[0,\infty)$ is measurable, non-negative, and non-decreasing, then an optimal solution π_{opt} to Problem (13) is obtained by Algorithm 1, where the function

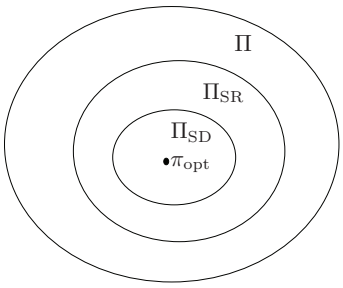

Fig. 5. Illustration of the set of causally feasible policies Π, the set of stationary randomized policies Π_{SR}, the set of stationary deterministic policies Π_{SD}, and the obtained optimal policy π_{opt}.

$z_\nu(\cdot)$ is determined by

$$z_\nu(y) = \sup\{z \in [0, M] : \mathbb{E}\left[g(y+z+Y')|Y = y\right] \leq \nu\}, \quad (17)$$

and (Y, Y') follow the same distribution as (Y_0, Y_1).

Proof Sketch of Theorem 3: We use Lagrangian duality theory to solve Problem (13). Different from traditional finite dimensional optimization problems [30], Problem (13) is an infinite dimensional functional optimization problem. Therefore, the Karush-Kuhn-Tucker (KKT) theorem for infinite dimensional space [31], [32] and the calculus of variations are required in the analysis. In particular, since the Lagrangian may not be strictly convex for some penalty functions, one-sided Gâteaux derivative (similar to sub-gradient in finite dimensional space) is used to solve the KKT conditions in Lebesgue space $L^2(\mu_Y)$. The proof details are provided in Appendix D. ∎

The policy spaces Π, Π_{SR}, Π_{SD}, and the obtained optimal policy π_{opt} are depicted in Fig. 5.

IV. When Is It Better to Wait than to Update?

When $\mathbb{E}[Y] \geq 1/f_{\max}$, the constraint (16) is always satisfied. In this case, a logical policy is the *zero-wait* policy: the source submits a fresh update once the prior update completes service, i.e., $\pi_{\text{zero-wait}} = (0, 0, \ldots)$. According to the example in the introduction, this zero-wait policy is not always optimal. When $\mathbb{E}[Y] < 1/f_{\max}$, due to the constraint (16), the minimum possible average waiting time is $\mathbb{E}[z(Y)] = 1/f_{\max} - \mathbb{E}[Y]$. However, even in this case, the optimal policy may have additional waiting time such that $\mathbb{E}[Y + z(Y)] > 1/f_{\max}$. In this section, we will study when it is optimal to submit updates with the minimum average waiting time and when it is not.

A. A Special Case of $g(\Delta) = \Delta$ with i.i.d. Service Times

Consider the case that $g(\Delta) = \Delta$ and the Y_i's are *i.i.d.* with $0 < \mathbb{E}[Y] < \infty$. In this case, Problem (13) has a simpler solution than that provided by Algorithm 1. Interestingly, this solution explicitly characterizes whether the optimal control $z(\cdot)$ can have the minimum average waiting time such that $\mathbb{E}[Y + z(Y)] = 1/f_{\max}$.

As shown in Fig. 3, $Q_n = \frac{1}{2}\left[(Y_n + Z_n + Y_{n+1})^2 - Y_n^2\right]$ is the area of a trapezoid. This corresponds to

$$q(y, z, y') = \frac{1}{2}\left[(y + z + y')^2 - y^2\right].$$

Algorithm 2 Bisection method for solving Problem (19)

given $l = 0$, sufficiently large u, tolerance ϵ.
repeat
  $\beta := (l + u)/2$.
  $o := \mathbb{E}\left[(Y + z(Y))\right] - \max\left(\frac{1}{f_{\max}}, \frac{\mathbb{E}\left[(Y+z(Y))^2\right]}{2\beta}\right)$,
  where $z(\cdot)$ is given by (21).
  if $o \geq 0$, $u := \beta$; **else**, $l := \beta$.
until $u - l \leq \epsilon$.
Compute $z(\cdot)$ by (21).
return $z(\cdot)$.

Because the Y_i's are *i.i.d.*, Y and Y' in Problem (13) are also *i.i.d.* Using this, we can obtain

$$\begin{aligned}
&\mathbb{E}[q(Y, z(Y), Y')] \\
=&\mathbb{E}\left[\frac{1}{2}(Y + z(Y) + Y')^2 - \frac{1}{2}Y'^2\right] \quad (18)\\
=&\frac{1}{2}\mathbb{E}\left[(Y + z(Y))^2\right] + \mathbb{E}\left[Y + z(Y)\right]\mathbb{E}\left[Y'\right],
\end{aligned}$$

where in (18) we have used that $\mathbb{E}[Y^2] = \mathbb{E}[Y'^2]$. Hence, Problem (13) can be reformulated as

$$\min_{z \in L^2(\mu_Y)} \frac{\mathbb{E}[(Y + z(Y))^2]}{2\mathbb{E}[Y + z(Y)]} + \mathbb{E}[Y] \quad (19)$$

$$\text{s.t.} \quad \mathbb{E}[Y + z(Y)] \geq \frac{1}{f_{\max}} \quad (20)$$
$$0 \leq z(y) \leq M, \ \forall\ y \geq 0.$$

The following lemma tells us that Problem (19) is a functional convex optimization problem.

Lemma 2: The functional $h_1 : L^2(\mu_Y) \to \mathbb{R}$ defined by

$$h_1(z) = \frac{\mathbb{E}[(Y + z(Y))^2]}{\mathbb{E}[Y + z(Y)]}$$

is convex on the domain

$$\textbf{dom}\ h_1 = \left\{z \in L^2(\mu_Y) : z(y) \in [0, M],\ \forall y \geq 0\right\}.$$

Proof: See Appendix E. ∎

Using the KKT theorem for infinite dimensional space and the calculus of variations, we can obtain

Theorem 4: If $\mathbb{E}[Y] > 0$, the optimal solution to Problem (19) is

$$z(y) = (\beta - y)_0^M, \quad (21)$$

where $(x)_0^M \triangleq \min\{\max\{x, 0\}, M\}$ and $\beta > 0$ satisfies

$$\mathbb{E}\left[Y + z(Y)\right] = \max\left(\frac{1}{f_{\max}}, \frac{\mathbb{E}[(Y + z(Y))^2]}{2\beta}\right). \quad (22)$$

Proof: See Appendix F. ∎

Equation (21) has the form of a water-filling solution, where the water-level β is given by the root of equation (22). One can observe that (17) reduces to (21) if $g(\Delta) = \Delta$, the Y_i's are *i.i.d.*, and ν is replaced by $\beta + \mathbb{E}[Y]$. The root β of equation (22) can be simply solved by the bisection search method in Algorithm 2. We note that Algorithm 2 has lower complexity than Algorithm 1 in the special case of $g(\Delta) = \Delta$ and *i.i.d.*

service process, while Algorithm 1 can obtain the optimal policy in more general scenarios.

Theorem 4 provides a closed-form criterion on whether the optimal $z(\cdot)$ satisfies $\mathbb{E}[Y + z(Y)] = 1/f_{\max}$. Specifically, (21) and (22) tell us that if $1/f_{\max} \geq \frac{\mathbb{E}[(Y+z(Y))^2]}{2\beta}$, then the optimal control $z(\cdot)$ satisfies

$$\mathbb{E}[Y + z(Y)] = \frac{1}{f_{\max}} \geq \frac{\mathbb{E}[(Y + z(Y))^2]}{2\beta}, \tag{23}$$

such that the optimal policy has the minimum possible average waiting time; otherwise, if $1/f_{\max} < \frac{\mathbb{E}[(Y+z(Y))^2]}{2\beta}$, the optimal control $z(\cdot)$ satisfies

$$\mathbb{E}[Y + z(Y)] = \frac{\mathbb{E}[(Y + z(Y))^2]}{2\beta} > \frac{1}{f_{\max}}, \tag{24}$$

such that the optimal policy can not achieve the minimum possible average waiting time.

Furthermore, if $\mathbb{E}[Y] \geq 1/f_{\max}$ (for example $f_{\max} = \infty$), the zero-wait policy is feasible. Meanwhile, the constraint (20) is always satisfied and can be removed. In this case, the optimality of the zero-wait policy is characterized by the following theorem.

Theorem 5: If $\mathbb{E}[Y] > 0$ and the zero-wait policy is feasible, then the zero-wait policy is optimal for Problem (19) if, and only if,

$$\frac{\mathbb{E}[Y^2]}{2\mathbb{E}[Y]} \leq y_{\inf}, \tag{25}$$

where $y_{\inf} = \inf\{y \in [0, \infty) : \Pr[Y \leq y] > 0\}$.

Proof: See Appendix G. ∎

Informally speaking, $y_{\inf}$ is the smallest possible value of the random service time Y. From Theorem 5, it immediately follows that:

Corollary 1: If $\mathbb{E}[Y] > 0$ and the zero-wait policy is feasible, then the following assertions are true:

(a). If the service times are positive and constant such that $Y = const > 0$, the zero-wait policy is optimal for Problem (19).

(b). If the service time distribution satisfies $y_{\inf} = 0$, the zero-wait policy is not optimal for Problem (19).

Remark: As one can readily see from Corollary 1(b), the zero-wait policy is not optimal for many commonly used service time distributions in communication and queueing theory, such as exponential distribution, geometric distribution, Erlang distribution, hyperexponential distribution, etc.

1) Simulation Results: We use "Optimal policy" to refer to the policy provided in Theorem 3 (or its special case in Theorem 4), and compare it with two reference policies:

- "Constant wait": Each update is followed by a constant waiting time $Z = 1/f_{\max} - \mathbb{E}[Y]$ before submitting the next update.
- "Minimum wait": The update waiting time is determined by $Z = z(Y)$, where $z(\cdot)$ is given by (21) and β in (22) is chosen to satisfy $\mathbb{E}[z(Y)] = 1/f_{\max} - \mathbb{E}[Y]$.[3]

When $\mathbb{E}[Y] = 1/f_{\max}$, both the constant wait and minimum wait policies reduce to the zero-wait policy. Two service process models are considered: The first is a discrete Markov chain with a probability mass function $\Pr[Y_i = 0] = \Pr[Y_i = 2] = 0.5$ and a transition matrix

$$P = \begin{bmatrix} p & 1-p \\ 1-p & p \end{bmatrix}.$$

Hence, the Y_i's are *i.i.d.* when $p = 0.5$, and the correlation coefficient between Y_i and Y_{i+1} is $\rho = 2p - 1$. The second is a log-normal distributed Markov chain, where $Y_i = e^{\sigma X_i}/\mathbb{E}[e^{\sigma X_i}]$ and $(X_0, X_1, \ldots)$ is a Gaussian Markov process satisfying the first-order AR equation

$$X_{i+1} = \eta X_i + \sqrt{1 - \eta^2} W_i,$$

where $\sigma > 0$ is the scale parameter of log-normal distribution, $\eta \in [-1, 1]$ is the parameter of the AR model, and the W_i's are *i.i.d.* Gaussian random variables with zero mean and unit variance. The log-normal distributed Markov chain is normalized such that $\mathbb{E}[Y_i] = 1$. According to the properties of log-normal distribution, the correlation coefficient between Y_i and Y_{i+1} is $\rho = (e^{\eta} - 1)/(e - 1)$. Then, the Y_i's are *i.i.d.* when $\eta = 0$.

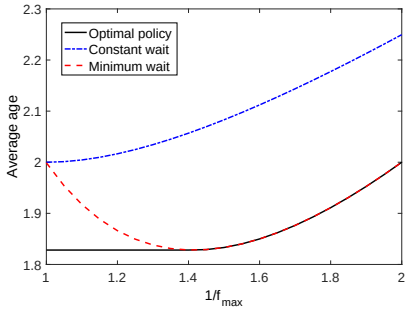

Fig. 6. Average age vs. $1/f_{\max}$ with *i.i.d.* discrete service times.

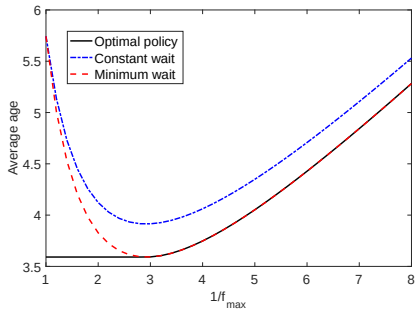

Fig. 7. Average age vs. $1/f_{\max}$ with *i.i.d.* log-normal distributed service times.

Figures 6 and 7 illustrate the average age vs. $1/f_{\max}$ for *i.i.d.* discrete and log-normal distributed service times, respectively. In both figures, one can observe that the constant wait policy always incurs a larger average age than the optimal policy. In addition, as expected from (23) and (24), as $1/f_{\max}$ exceeds a certain threshold, the optimal policy meets the constraint (20) with equality. For smaller values of $1/f_{\max}$, the constraint (20) is not active in the optimal solution. Consequently, the minimum wait policy deviates from the optimal policy for small values of $1/f_{\max}$, which is in accordance to Corollary 1(b).

[3]This policy was called "β-minimum" in [11].

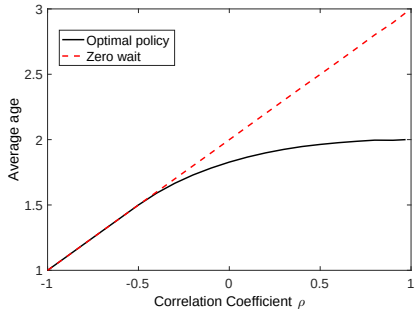

Fig. 8. Average age vs. the correlation coefficient ρ between Y_i and Y_{i+1} for discrete service times, where $\mathbb{E}[Y] \geq 1/f_{\max}$.

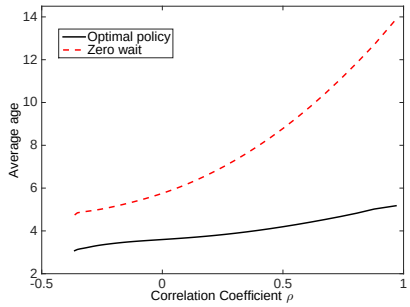

Fig. 9. Average age vs. the correlation coefficient ρ between Y_i and Y_{i+1} for log-normal distributed service times, where $\mathbb{E}[Y] \geq 1/f_{\max}$.

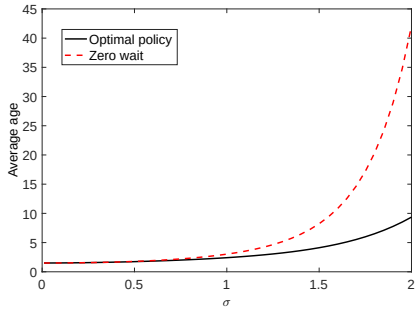

Fig. 10. Average age vs. the distribution parameter σ of $i.i.d.$ log-normal distributed service times, where $\mathbb{E}[Y] \geq 1/f_{\max}$.

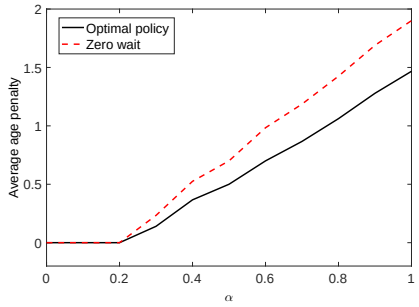

Fig. 11. Average age penalty vs. the parameter α of stair-step penalty functions with discrete service times, where $\mathbb{E}[Y] \geq 1/f_{\max}$, $g(\Delta) = \lfloor \alpha\Delta \rfloor$, and $\rho = 0.4$.

B. General Age Penalties and Correlated Service Processes

For general age penalties and correlated service processes, it is essentially difficult to find closed-form characterization on whether the optimal control $z(\cdot)$ can have the minimum average waiting time such that $\mathbb{E}[Y + z(Y)] = 1/f_{\max}$. Therefore, we focus on the case of $\mathbb{E}[Y] \geq 1/f_{\max}$ (this is equivalent to removing the update frequency constraint (16)) and study when the zero-wait policy minimizes the average age penalty. Sufficient conditions for the optimality of the zero-wait policy are provided as follows:

Lemma 3: Suppose that $\mathbb{E}[Y] \geq 1/f_{\max}$, $g(\cdot)$ is measurable, non-negative, and non-decreasing. The zero-wait policy is optimal for Problem (13) if one of the following conditions is satisfied:

1). The correlation coefficient between Y_i and Y_{i+1} is -1;
2). The Y_i's are equal to a constant value;
3). $g(\cdot)$ is a constant function.

Proof: See Appendix H. ∎

1) Simulation Results: We now provide some simulation results for general age penalties and/or correlated service processes. Figures 8 and 9 depict the average age vs. the correlation coefficient ρ between Y_i and Y_{i+1} for discrete and log-normal distributed service times, respectively. In Fig. 8, the range of ρ is $[-1, 1)$. We observe that the zero-wait policy is optimal when $\rho \in [-1, -0.5]$, and the performance gap between the optimal policy and the zero-wait policy grows with ρ when $\rho \geq -0.5$. This is in accordance with the example in the introduction: As ρ grows, the occurrence of two consecutive zero service times (i.e., $(Y_i, Y_{i+1}) = (0, 0)$) increases. Therefore, more and more updates are wasted in the zero-wait policy, leading to a larger gap from the optimum. In Fig. 9, the range of ρ is $[(e^{-1} - 1)/(e - 1), 1)$ and the sub-optimality gap of the zero-wait policy also increases with ρ. The point $\rho = 1$ is not plotted in these figures because the corresponding Markov chains are not ergodic.

Figure 10 considers the average age vs. the parameter σ of log-normal distributed service times, where $\rho = (e^{0.5} - 1)/(e - 1)$. We observe that the zero-wait policy is optimal for small σ and is not optimal for large σ. When $\sigma = 0$, the service times are constant, i.e., $Y_i = 1$ for all i, and hence by Lemma 3, the zero-wait policy is optimal. For large σ, the time-average age of the zero-wait policy is significantly larger than the optimum. This implies that the sub-optimality gap of the zero-wait policy can be quite large for heavy-tail service time distributions.

Figures 11-16 show the average age penalty vs. the parameter α of three types of age penalty functions, where the stair-step function $g(\Delta) = \lfloor \alpha\Delta \rfloor$ is considered in Fig. 11 and 12, the power function $g(\Delta) = \Delta^\alpha$ is considered in Fig. 13 and 14, and the exponential function $g(\Delta) = e^{\alpha\Delta} - 1$ is considered in Fig. 15 and 16. The correlation coefficient is $\rho = 0.4$ for discrete service times, and is $\rho = (e^{0.5} - 1)/(e - 1)$ for log-normal distributed service times. We find that the zero-wait policy is optimal if $\alpha = 0$, in which case $g(\Delta)$ is a constant function. When $\alpha > 0$, the zero-wait policy may not be optimal.

These simulation results suggest that the conditions in Lemma 3 are sufficient but not necessary.

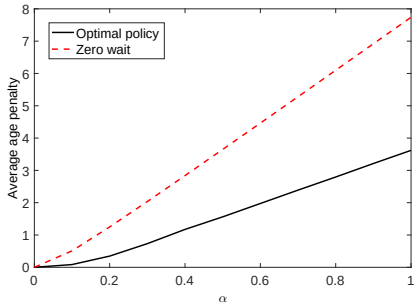

Fig. 12. Average age penalty vs. the parameter α of stair-step penalty functions with log-normal distributed service times, where $\mathbb{E}[Y] \geq 1/f_{\max}$, $g(\Delta) = \lfloor \alpha \Delta \rfloor$, and $\rho = (e^{0.5} - 1)/(e - 1)$.

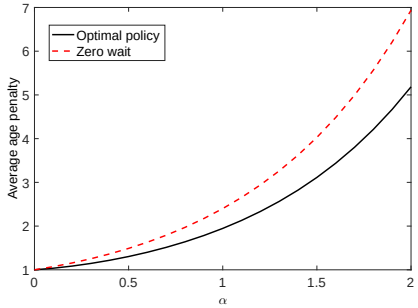

Fig. 13. Average age penalty vs. the parameter α of power penalty functions with discrete service times, where $\mathbb{E}[Y] \geq 1/f_{\max}$, $g(\Delta) = \Delta^{\alpha}$, and $\rho = 0.4$.

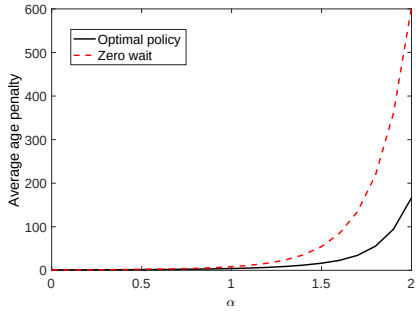

Fig. 14. Average age penalty vs. the parameter α of power penalty functions with log-normal distributed service times, where $1/f_{\max} \leq \mathbb{E}[Y]$, $g(\Delta) = \Delta^{\alpha}$, and $\rho = (e^{0.5} - 1)/(e - 1)$.

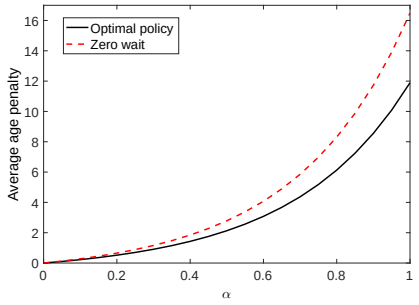

Fig. 15. Average age penalty vs. the parameter α of exponential penalty functions with discrete service times, where $1/f_{\max} \leq \mathbb{E}[Y]$, $g(\Delta) = e^{\alpha \Delta} - 1$, and $\rho = 0.4$.

V. Conclusion

We studied the optimal control of status updates from a source to a remote destination via an FCFS communication server. We generalized the existing studies on the age-of-information to cover general age penalty functions and non-*i.i.d.* service processes. We developed efficient algorithms to find the optimal status update policy for minimizing the average age penalty. We showed that, surprisingly, in many scenarios, the optimal policy is to wait for a certain amount of time before submitting a new update. In particular, the widely-adopted zero-wait policy can be far from the optimum if (i) the penalty function grows quickly with respect to the age, and (ii) the update service times are positively correlated and have heavy tail distributions.

Appendix A
Proof of Theorem 1

A. An Upper Bound of $\overline{g}_{opt}$

By restricting Π in Problem (8) to Π_{SR}, we obtain the following problem:

$$\overline{g}_{\text{SR}} = \min_{\pi \in \Pi_{\text{SR}}} \limsup_{n \to \infty} \frac{\mathbb{E}\left[\sum_{i=0}^{n-1} q(Y_i, Z_i, Y_{i+1})\right]}{\mathbb{E}[\sum_{i=0}^{n-1} (Y_i + Z_i)]} \quad \text{s.t.} \quad \liminf_{n \to \infty} \frac{1}{n} \mathbb{E}\left[\sum_{i=0}^{n-1} (Y_i + Z_i)\right] \geq 1/f_{\max}, \tag{26}$$

where $\overline{g}_{\text{SR}}$ is the optimum objective value of Problem (26). Since $\Pi_{\text{SR}} \subseteq \Pi$, we can obtain

$$\overline{g}_{\text{SR}} \geq \overline{g}_{\text{opt}}. \tag{27}$$

It is easy to show that the (Y_i, Z_i, Y_{i+1})'s are stationary and ergodic for all stationary randomized policies. This, together with the condition that $g(\cdot)$ is measurable, tells us that $q(Y_i, Z_i, Y_{i+1})$ is stationary and ergodic [33, Theorems 7.1.1 and 7.1.3]. For any stationary randomized policy $\pi = (Z_0, Z_1, \ldots) \in \Pi_{\text{SR}}$, we obtain

$$\frac{1}{n} \mathbb{E}\left[\sum_{i=0}^{n-1} q(Y_i, Z_i, Y_{i+1})\right] = \mathbb{E}[q(Y_0, Z_0, Y_1)], \tag{28}$$

$$\frac{1}{n} \mathbb{E}\left[\sum_{i=0}^{n-1} (Y_i + Z_i)\right] = \mathbb{E}[Y_0 + Z_0]. \tag{29}$$

Hence, Problem (26) can be reformulated as Problem (12).

B. The Upper Bound of $\overline{g}_{opt}$ is Tight, i.e., $\overline{g}_{SR} = \overline{g}_{opt}$

We will show $\overline{g}_{\text{SR}} = \overline{g}_{\text{opt}}$ in 4 steps. The following definitions are needed: Since $\overline{g}_{\text{opt}}$ is finite, for each causally feasible policy $\pi = (Z_0, Z_1, \ldots) \in \Pi$ we can define $a_{n,\pi}$ and $b_{n,\pi}$ as in (10) and (11), respectively.

Further, define Γ_{SR} as the set of limit points of sequences $\big((a_{n,\pi}, b_{n,\pi}), n = 1, 2, \ldots\big)$ associated with all stationary randomized policies $\pi \in \Pi_{\text{SR}}$. Because the reward $q(Y_i, Z_i, Y_{i+1})$ and interval $Y_i + Z_i$ are stationary and ergodic for all stationary randomized policies $\pi \in \Pi_{\text{SR}}$, the sequence $(a_{n,\pi}, b_{n,\pi})$ has a

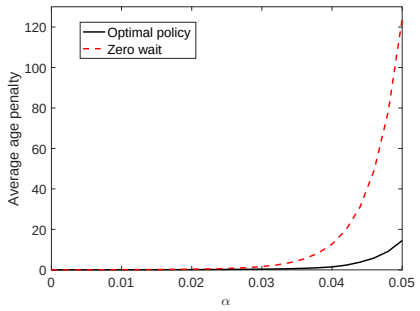

Fig. 16. Average age penalty vs. the parameter α of exponential penalty functions with log-normal distributed service times, where $1/f_{\max} \leq \mathbb{E}[Y]$, $g(\Delta) = e^{\alpha\Delta} - 1$, and $\rho = (e^{0.5} - 1)/(e - 1)$.

unique limit point in the form of

$$\left(\mathbb{E}[q(Y, Z, Y')] - \overline{g}_{\text{opt}}\mathbb{E}[Y + Z], \mathbb{E}[Y + Z]\right). \tag{30}$$

Hence, Γ_{SR} is the set of all points $(\mathbb{E}[q(Y, Z, Y')] - \overline{g}_{\text{opt}}\mathbb{E}[Y + Z], \mathbb{E}[Y + Z])$, where each point is associated with a conditional probability measure $p(y, A) = \Pr[Z \in A|Y = y]$, and the measure of (Y, Y') is the same as that of (Y_0, Y_1). Note that $\mathbb{E}[Y] = \mathbb{E}[Y']$.

Step 1: We will show that Γ_{SR} is a convex and compact set.

It is easy to show that Γ_{SR} is convex by considering a stationary randomized policy that is a mixture of two stationary randomized policies.

For compactness, let $((d_j, e_j), j = 1, 2, \cdots)$ be any sequence of points in Γ_{SR}, we need to show that there is a convergent subsequence (d_{j_k}, e_{j_k}) whose limit is also in Γ_{SR}. Since $(d_j, e_j) \in \Gamma_{\text{SR}}$, there must exist $(Y, Z_{(j)}, Y')$ with conditional probability $p_j(y, A) = \Pr[Z_{(j)} \in A|Y = y]$, such that $d_j = \mathbb{E}[q(Y, Z_{(j)}, Y')] - \overline{g}_{\text{opt}}\mathbb{E}[Y + Z_{(j)}]$, $e_j = \mathbb{E}[Y + Z_{(j)}]$. Let μ_j be the joint probability measure of $(Y, Z_{(j)}, Y')$, then (d_j, e_j) is uniquely determined by μ_j. For any L satisfying $L \geq M$, we can obtain

$$\begin{aligned}
&\mu_j(Y \leq L, Z_{(j)} \leq L, Y' \leq L) \\
=& \Pr(Y \leq L, Y' \leq L) \\
\geq& \Pr(Y + Y' \leq L) \\
\geq& 1 - \frac{\mathbb{E}[Y + Y']}{L}, \ \forall\ j,
\end{aligned}$$

where the equality is due to the fact that $Z_{(j)} \leq M \leq L$ and the last inequality is due to Markov's inequality. Therefore, for any ϵ, there is an L such that

$$\liminf_{j\to\infty} \mu_j(|Y| \leq L, |Z_{(j)}| \leq L, |Y'| \leq L) \geq 1 - \epsilon.$$

Hence, the sequence of measures μ_j is tight. By Helly's selection theorem [33, Theorem 3.9.2], there is a subsequence of measures μ_{j_k} that converges weakly to a limit measure μ_∞.

Let $(Y, Z_{(\infty)}, Y')$ and $p_\infty(y, A) = \Pr[Z_\infty \in A|Y = y]$ denote the random vector and conditional probability corresponding to the limit measure μ_∞, respectively. We can define $d_\infty = \mathbb{E}[q(Y, Z_{(\infty)}, Y')] - \overline{g}_{\text{opt}}\mathbb{E}[Y + Z_{(\infty)}]$, $e_\infty = \mathbb{E}[Y + Z_{(\infty)}]$. Since the function $q(y, z, y')$ is in the form of an integral, it is continuous and thus measurable. Using the continuous mapping theorem [33, Theorem 3.2.4], we can obtain that $q(Y, Z_{(j_k)}, Y')$ converges weakly to $q(Y, Z_{(\infty)}, Y')$. Then, using the condition (6), together with the dominated convergence theorem (Theorem 1.6.7 of [33]) and Theorem 3.2.2 of [33], we can obtain $\lim_{k\to\infty}(d_{j_k}, e_{j_k}) = (d_\infty, e_\infty)$. Hence, $((d_j, e_j), j = 1, 2, \cdots)$ has a convergent subsequence. Further, we can generate a stationary randomized policy $\pi_{\infty,\text{SR}}$ by using the conditional probability $p_\infty(y, A)$ corresponding to μ_∞. Then, (d_∞, e_∞) is the limit point generated by the stationary randomized policy $\pi_{\infty,\text{SR}}$, which implies $(d_\infty, e_\infty) \in \Gamma_{\text{SR}}$. In summary, any sequence (d_j, e_j) in Γ_{SR} has a convergent subsequence (d_{j_k}, e_{j_k}) whose limit (d_∞, e_∞) is also in Γ_{SR}. Therefore, Γ_{SR} is a compact set.

Step 2: We will show that there exists an optimal policy $\pi_{opt} \in \Pi$ of Problem (8) such that the sequence $(a_{n,\pi_{opt}}, b_{n,\pi_{opt}})$ associated with policy π_{opt} has at least one limit point in Γ_{SR}.

Since the sequence $(Y_0, Y_1, \ldots)$ is a Markov chain, the observation Y_{i+1} depends only on the immediately preceding state Y_i and not on the history state and control $Y_0, \ldots, Y_{i-1}$, $Z_0, \ldots, Z_{i-1}$. Therefore, Y_i is the *sufficient statistic* [17, p. 252] for solving Problem (8). This tells us that there exists an optimal policy $\pi_{\text{opt}} = (Z_0, Z_1, \ldots) \in \Pi$ of Problem (8) in which the control action Z_i is determined based on only Y_i, but not the history state and control $Y_0, \ldots, Y_{i-1}$, $Z_0, \ldots, Z_{i-1}$ [17]. We will show that the sequence $(a_{n,\pi_{\text{opt}}}, b_{n,\pi_{\text{opt}}})$ associated with this policy π_{opt} has at least one limit point in Γ_{SR}.

It is known that Z_i takes values in the standard Borel space $(\mathbb{R}, \mathcal{R})$, where $\mathcal{R}$ is the Borel σ-field. According to [33, Thoerem 5.1.9], for each i there exists a conditional probability measure $p'_i(y, A)$ such that $p'_i(y, A) = \Pr(Z_i \in A|Y_i = y)$ for almost all y. That is, the control action Z_i is determined based on Y_i and the conditional probability measure $p'_i(y, A) = \Pr(Z_i \in A|Y_i = y)$. One can use this conditional probability $p'_i(y, A)$ to generate a stationary randomized policy $\pi'_{i,\text{SR}} \in \Pi_{\text{SR}}$. Then, the one-stage expectation $(\mathbb{E}[q(Y_i, Z_i, Y_{i+1})] - \overline{g}_{\text{opt}}\mathbb{E}[Y_i + Z_i], \mathbb{E}[Y_i + Z_i])$ is exactly the limit point generated by the stationary randomized policy $\pi'_{i,\text{SR}}$. Thus, $(\mathbb{E}[q(Y_i, Z_i, Y_{i+1})] - \overline{g}_{\text{opt}}\mathbb{E}[Y_i + Z_i], \mathbb{E}[Y_i + Z_i]) \in \Gamma_{\text{SR}}$ for all $i = 0, 1, 2, \ldots$ Using (10), (11), and the fact that Γ_{SR} is convex, we can obtain $(a_{n,\pi_{\text{opt}}}, b_{n,\pi_{\text{opt}}}) \in \Gamma_{\text{SR}}$ for all $n = 1, 2, 3 \ldots$ In other words, the sequence $(a_{n,\pi_{\text{opt}}}, b_{n,\pi_{\text{opt}}})$ is within Γ_{SR}. Since Γ_{SR} is a compact set, the sequence $(a_{n,\pi_{\text{opt}}}, b_{n,\pi_{\text{opt}}})$ must have a convergent subsequence, whose limit is in Γ_{SR}.

Step 3: Let $(a^*, b^*) \in \Gamma_{SR}$ be one limit point of the sequence $(a_{n,\pi_{opt}}, b_{n,\pi_{opt}})$ associated with policy π_{opt}. We will show that $a^* \leq 0$ and $b^* \geq 1/f_{\max}$.

Policy π_{opt} is feasible for Problem (8) and meanwhile achieves the optimum objective value $\overline{g}_{\text{opt}}$. Hence,

$$\limsup_{n\to\infty} \frac{c_{n,\pi_{\text{opt}}}}{b_{n,\pi_{\text{opt}}}} = \overline{g}_{\text{opt}}, \tag{31}$$

$$\liminf_{n\to\infty} b_{n,\pi_{\text{opt}}} \geq 1/f_{\max}, \tag{32}$$

where

$$c_{n,\pi_{\text{opt}}} \triangleq \frac{1}{n}\mathbb{E}\left[\sum_{i=0}^{n-1} q(Y_i, Z_i, Y_{i+1})\right].$$

By (11), $b_{n,\pi_{\text{opt}}}$ is upper bounded by

$$b_{n,\pi_{\text{opt}}} \leq M + \mathbb{E}[Y] < \infty.$$

Hence, by (10), we have

$$\begin{aligned} a_{n,\pi_{\text{opt}}} &= c_{n,\pi_{\text{opt}}} - \overline{g}_{\text{opt}} b_{n,\pi_{\text{opt}}} \\ &\leq \max\{c_{n,\pi_{\text{opt}}} - \overline{g}_{\text{opt}} b_{n,\pi_{\text{opt}}}, 0\} \\ &= \max\{\frac{c_{n,\pi_{\text{opt}}}}{b_{n,\pi_{\text{opt}}}} - \overline{g}_{\text{opt}}, 0\} b_{n,\pi_{\text{opt}}} \\ &\leq \max\{\frac{c_{n,\pi_{\text{opt}}}}{b_{n,\pi_{\text{opt}}}} - \overline{g}_{\text{opt}}, 0\}(M + \mathbb{E}[Y]). \end{aligned}$$

Taking the $\limsup$ in this inequality and using (31), yields

$$\limsup_{n\to\infty} a_{n,\pi_{\text{opt}}} \leq 0. \tag{33}$$

Because (a^*, b^*) is one limit point of $(a_{n,\pi_{\text{opt}}}, b_{n,\pi_{\text{opt}}})$, we have

$$a^* \leq \limsup_{n\to\infty} a_{n,\pi_{\text{opt}}}, b^* \geq \liminf_{n\to\infty} b_{n,\pi_{\text{opt}}}. \tag{34}$$

By (32)-(34), we have $a^* \leq 0$ and $b^* \geq 1/f_{\max}$.

Step 4: We will show that there exists a stationary randomized policy that is optimal for Problems (8) *and* (12)*, and thus* $\overline{g}_{SR} = \overline{g}_{opt}$. By the definition of Γ_{SR}, $(a^*, b^*) \in \Gamma_{\text{SR}}$ must be the limit point generated by a stationary randomized policy $\pi^* \in \Pi_{\text{SR}}$. Let (Y, Z^*, Y') be a random vector with the stationary distribution of policy π^*. Then, (30) implies

$$(a^*, b^*) = \left(\mathbb{E}[q(Y, Z^*, Y')] - \overline{g}_{\text{opt}}\mathbb{E}[Y + Z^*], \mathbb{E}[Y + Z^*]\right).$$

Using $a^* \leq 0$ and $b^* \geq 1/f_{\max}$, we can obtain

$$\mathbb{E}[q(Y, Z^*, Y')] - \mathbb{E}[Y + Z^*]\overline{g}_{\text{opt}} \leq 0, \tag{35}$$

$$\mathbb{E}[Y + Z^*] \geq 1/f_{\max}. \tag{36}$$

By (35) and $\mathbb{E}[Y + Z^*] > 0$, we have

$$\frac{\mathbb{E}[q(Y, Z^*, Y')]}{\mathbb{E}[Y + Z^*]} \leq \overline{g}_{\text{opt}}.$$

Further, the inequality (36) suggests that the stationary randomized policy π^* is feasible for Problem (12). Hence,

$$\frac{\mathbb{E}[q(Y, Z^*, Y')]}{\mathbb{E}[Y + Z^*]} \geq \overline{g}_{\text{SR}}.$$

Therefore, $\overline{g}_{\text{SR}} \leq \overline{g}_{\text{opt}}$. This and (27) suggest that

$$\frac{\mathbb{E}[q(Y, Z^*, Y')]}{\mathbb{E}[Y + Z^*]} = \overline{g}_{\text{SR}} = \overline{g}_{\text{opt}}.$$

This completes the proof.

Appendix B
Proof of Theorem 2

Consider an arbitrarily chosen stationary randomized policy $\pi_1 \in \Pi_{\text{SR}}$ that is feasible for problem (12). We will show that there exists a feasible stationary deterministic policy that is no worse than policy π_1.

For any y, we can use the conditional probability $p(y, A)$ associated with policy π_1 to compute the conditional expectation $\mathbb{E}[Z|Y = y]$ by

$$\mathbb{E}[Z|Y = y] = \int_0^M z p(y, dz).$$

Since the conditional expectation $\mathbb{E}[Z|Y]$ is unique w.p.1 [33, Section 5.1], there is a deterministic function $z(\cdot)$ such that $z(y) = \mathbb{E}[Z|Y = y]$ w.p.1. Consider the set $\Lambda \subset \Pi_{\text{SR}}$ of all stationary randomized policies that satisfy $\mathbb{E}[Z|Y = y] = z(y)$ w.p.1. Then, the stationary randomized policy π_1 is in Λ. It is also easy to show that the stationary deterministic policy $(Z_i = z(Y_i), i = 1, 2, \ldots)$ is also in Λ.

Using the iterated expectation, for any policy in Λ

$$\mathbb{E}[Y + Z] = \mathbb{E}\big[Y + \mathbb{E}[Z|Y]\big] = \mathbb{E}\big[Y + z(Y)\big]. \tag{37}$$

Because $\pi_1 \in \Lambda$ is feasible for problem (12), any policy in Λ is feasible for problem (12).

Since $q(y, z, y')$ is the integral of a non-decreasing function g, it is easy to show that the function $q(y, \cdot, y')$ is convex. For any policy $\pi \in \Lambda$, Jensen's inequality tells us that

$$\begin{aligned} &\mathbb{E}[q(Y, Z, Y')|Y, Y'] \\ &\geq q(Y, \mathbb{E}[Z|Y, Y'], Y') \\ &= q(Y, \mathbb{E}[Z|Y], Y') \\ &= q(Y, z(Y), Y'), \text{ (w.p.1)}, \end{aligned} \tag{38}$$

where (38) is due to the fact that Z is determined based on Y, but not Y'. Taking the expectation over (Y, Y'), yields

$$\mathbb{E}[q(Y, z(Y), Y')] \leq \mathbb{E}[q(Y, Z, Y')]$$

for any policy $\pi \in \Lambda$, where equality holds if $Z = z(Y)$. This and (37) suggest that the stationary deterministic policy $(Z_i = z(Y_i), i = 1, 2, \ldots)$ achieves the smallest objective value for problem (12) among all policies in Λ. In conclusion, for any feasible stationary randomized policy $\pi_1 \in \Pi_{\text{SR}}$, we can find a feasible stationary deterministic policy that is no worse than policy π_1. This completes the proof.

Appendix C
Proof of Lemma 1

We need the following lemma:

Lemma 4: If $l : \mathbb{R} \to \mathbb{R}$ is a convex function, then the functional $w : L^2(\mu_Y) \to \mathbb{R}$ defined by

$$w(z) = \int_0^\infty l(z(y)) d\mu_Y(y) \tag{39}$$

is also convex.

Proof: For any $\lambda \in [0, 1]$ and $z_1, z_2 \in L^2(\mu_Y)$, we have

$$\begin{aligned} &w(\lambda z_1 + (1-\lambda) z_2) \\ =&\int_0^\infty l(\lambda z_1(y) + (1-\lambda) z_2(y)) d\mu_Y(y) \\ \leq&\int_0^\infty \left[\lambda l(z_1(y)) + (1-\lambda) l(z_2(y))\right] d\mu_Y(y) \\ =&\lambda w(z_1) + (1-\lambda) w(z_2). \end{aligned} \tag{40}$$

By this, $w(z)$ is convex. ■

We now prove Lemma 1. Since $q(y, z, y')$ is the integral of a non-negative and non-decreasing function g, it is easy to

show that the function $q(y,\cdot,y')$ is non-negative and convex. Hence, the conditional expectation $\mathbb{E}\left[q(y,\cdot,Y')|Y=y\right]$ is also convex. We can obtain

$$\begin{aligned}&\mathbb{E}\left[q(Y,z(Y),Y')\right]\\&=\int_0^\infty \mathbb{E}\left[q(y,z(y),Y')|Y=y\right]d\mu_Y(y).\end{aligned}\tag{41}$$

According to Lemma 4, $\mathbb{E}\left[q(Y,z(Y),Y')\right]$ is a convex functional of z and $\mathbb{E}[Y+z(Y)]$ is an affine functional of z. It is known that the ratio of a non-negative convex functional and positive affine functional is quasi-convex [30, p. 103]. Hence, $h(z)$ is quasi-convex, which completes the proof.

Appendix D
Proof of Theorem 3

We use the Lagrangian duality approach to solve Problem (15). The Lagrangian of Problem (15) is

$$\begin{aligned}&L(z,\zeta,\gamma,\tau)\\&=\int_0^\infty \mathbb{E}\left[q(y,z(y),Y')|Y=y\right]d\mu_Y(y)\\&\quad-c\int_0^\infty [y+z(y)]d\mu_Y(y)\\&\quad+\zeta\left[1/f_{\max}-\int_0^\infty [y+z(y)]d\mu_Y(y)\right]\\&\quad-\int_0^\infty \gamma(y)z(y)d\mu_Y(y)+\int_0^\infty \tau(y)(z(y)-M)d\mu_Y(y)\\&=\int_0^\infty\Big\{\mathbb{E}\left[q(y,z(y),Y')|Y=y\right]-(c+\zeta)[y+z(y)]\\&\quad-\gamma(y)z(y)+\tau(y)\Big[z(y)-M\Big]\Big\}d\mu_Y(y)+\zeta 1/f_{\max}.\end{aligned}\tag{42}$$

Since Problem (15) is feasible, all constraints are affine, the refined Slater's condition [30, Sec. 5.2.3] is satisfied. According to [31, Proposition 3.3.2] and [32, pp. 70-72], the Karush-Kuhn-Tucker (KKT) theorem remains valid for the Lebesgue space $L^2(\mu_Y)$. Hence, if a vector (z,ζ,γ,τ) satisfies the KKT conditions (43)-(49), it is an optimal solution to (15). The KKT conditions are given by

$$z=\min_{x\in L^2(\mu_Y)} L(x,\zeta,\gamma,\tau),\tag{43}$$

$$\zeta\geq 0,\int_0^\infty [y+z(y)]d\mu_Y(y)\geq 1/f_{\max},\tag{44}$$

$$\gamma(y)\geq 0,z(y)\geq 0,\forall\ y\geq 0,\tag{45}$$

$$\tau(y)\geq 0,z(y)\leq M,\forall\ y\geq 0,\tag{46}$$

$$\zeta\left[1/f_{\max}-\int_0^\infty [y+z(y)]d\mu_Y(y)\right]=0,\tag{47}$$

$$\gamma(y)z(y)=0,\forall\ y\geq 0,\tag{48}$$

$$\tau(y)(z(y)-M)=0,\forall\ y\geq 0.\tag{49}$$

We now solve the KKT conditions (43)-(49) by using the calculus of variations. The one-sided Gâteaux derivative (similar to sub-gradient in finite dimensional space) of a functional h in the direction of $w\in L^2(\mu_Y)$ at $z\in L^2(\mu_Y)$ is defined as

$$\delta h(z;w)\triangleq\lim_{\epsilon\to 0^+}\frac{h(z+\epsilon w)-h(z)}{\epsilon}.\tag{50}$$

If h is a function on $\mathbb{R}$, then (50) becomes the common one-sided derivative. Let $l(z,y,\zeta,\gamma,\tau)$ denote the integrand in (42), and $r(z,y)=\mathbb{E}\left[q(y,z(y),Y')|Y=y\right]$. According to Lemma 4, the function $q(y,z,y')$ and functionals $r(z,y)$, $l(z,y,\zeta,\gamma,\tau)$, and $L(z,\zeta,\gamma,\tau)$ are all convex in z. Therefore, their one-sided Gâteaux derivatives with respect to z exist [34, p. 709]. Since $g(x)$ is right-continuous, for any given (y,y'), the one-sided derivative $\delta q(y,z;w,y')$ of function $q(y,z,y')$ with respect to z is given by

$$\begin{aligned}&\delta q(y,z;w,y')\\&=\lim_{\epsilon\to 0^+}\frac{1}{\epsilon}\big\{q(y,z+\epsilon w,y')-q(y,z,y')\big\}\\&=\begin{cases}\lim\limits_{x\to z^+} g(y+x+y')w, & \text{if } w\geq 0;\\ \lim\limits_{x\to z^-} g(y+x+y')w, & \text{if } w<0.\end{cases}\end{aligned}$$

Next, consider the one-sided Gâteaux derivative $\delta r(z;w,y)$ of functional $r(z,y)$. Since the function $g:[0,\infty)\to[0,\infty)$ is non-decreasing, $z\to q(y,z,y')$ is convex and finite for all $z\in[0,M]$. Hence, the function $\epsilon\to[q(y,z+\epsilon w,y')-q(y,z,y')]/\epsilon$ is non-decreasing and bounded from above on $(0,1]$ [31, Proposition 1.1.2(i)]. By using the monotone convergence theorem [33, Theorem 1.5.6], we can interchange the limit and integral operators in $\delta r(z;w,y)$ such that

$$\begin{aligned}&\delta r(z;w,y)\\&=\lim_{\epsilon\to 0^+}\frac{1}{\epsilon}\mathbb{E}\left[q(y,z(y)+\epsilon w(y),Y')-q(y,z(y),Y')|Y=y\right]\\&=\mathbb{E}\left[\lim_{\epsilon\to 0^+}\frac{1}{\epsilon}\big\{q(y,z(y)+\epsilon w(y),Y')-q(y,z(y),Y')\big\}\middle|Y=y\right]\\&=\mathbb{E}\Bigg[\lim_{x\to z(y)^+} g(y+x+Y')w(y)1_{\{w(y)>0\}}\\&\quad+\lim_{x\to z(y)^-} g(y+x+Y')w(y)1_{\{w(y)<0\}}\middle|Y=y\Bigg]\\&=\lim_{x\to z(y)^+}\mathbb{E}\left[g(y+x+Y')w(y)1_{\{w(y)>0\}}\middle|Y=y\right]\\&\quad+\lim_{x\to z(y)^-}\mathbb{E}\left[g(y+x+Y')w(y)1_{\{w(y)<0\}}\middle|Y=y\right],\end{aligned}\tag{51}$$

where 1_E is the indicator function of event E. By using the monotone convergence theorem again, we have

$$\begin{aligned}&\delta L(z;w,\zeta,\gamma,\tau)\\&=\int_0^\infty \delta l(z;w,y,\zeta,\gamma,\tau)d\mu_Y(y)\\&=\int_0^\infty \delta r(z;w,y)d\mu_Y(y)\\&\quad+\int_0^\infty\left[-(c+\zeta)-\gamma(y)+\tau(y)\right]w(y)d\mu_Y(y).\end{aligned}\tag{52}$$

According to [34, p. 710], z is an optimal solution to (43) if and only if

$$\delta L(z;w,\zeta,\gamma,\tau)\geq 0,\quad\forall\ w\in L^2(\mu_Y).\tag{53}$$

Since $w(\cdot)$ is an arbitrary function in $L^2(\mu_Y)$, considering positive functions $w(y) > 0$, we can obtain from (51)-(53) that for each $y \geq 0$, $z(y)$ must satisfy

$$\lim_{x \to z(y)^+} \mathbb{E}\left[g(y+x+Y')|Y=y\right] - (c+\zeta) - \gamma(y) + \tau(y) \geq 0. \tag{54}$$

Similarly, considering negative functions $w(y) < 0$, we can obtain that for each $y \geq 0$, $z(y)$ must satisfy

$$\lim_{x \to z(y)^-} \mathbb{E}\left[g(y+x+Y')|Y=y\right] - (c+\zeta) - \gamma(y) + \tau(y) \leq 0. \tag{55}$$

Because $g(\cdot)$ is non-decreasing, we can obtain from (54) and (55) that for each $y \geq 0$, $z(y)$ needs to satisfy

$$\mathbb{E}\left[g(y+x+Y')|Y=y\right] - (c+\zeta) - \gamma(y) + \tau(y) \geq 0 \tag{56}$$

for all $x > z(y)$, and

$$\mathbb{E}\left[g(y+x+Y')|Y=y\right] - (c+\zeta) - \gamma(y) + \tau(y) \leq 0 \tag{57}$$

for all $x < z(y)$.

We solve the optimal primal solution $z(\cdot)$ by considering the following three cases:

Case 1: $\gamma(y) = \tau(y) = 0$. The solutions to (56) and (57) may not be unique, if function g is not strictly increasing. In particular, there may exist an interval $[a(y), b(y)]$ such that each $z(y) \in [a(y), b(y)]$ satisfies (56) and (57) for some y. In this case, we choose the largest possible solution of $z(y)$ to make sure that the constraint (14) is satisfied. The largest solution satisfying (56) and (57) is given by

$$z(y) = \sup\{x \in [0, M] : \mathbb{E}\left[g(y+x+Y')|Y=y\right] \leq c+\zeta\}, \quad \forall\ y \geq 0,$$

which is exactly (17).

Case 2: $\gamma(y) > 0$. By (48), we have $z(y) = 0$.

Case 3: $\tau(y) > 0$. By (49), we have $z(y) = M$.

In summary, the optimal primal solution $z(\cdot)$ is given by (17).

Next, we find the optimal dual variable ζ. By (44) and (47), the optimal ζ satisfies

$$\zeta = 0, \mathbb{E}\left[Y + z(Y)\right] \geq 1/f_{\max} \tag{58}$$

or

$$\zeta > 0, \mathbb{E}\left[Y + z(Y)\right] = 1/f_{\max}, \tag{59}$$

where $\mathbb{E}\left[Y + z(Y)\right]$ is determined by the optimal primal solution (17). Since $\mathbb{E}\left[Y + z(Y)\right]$ is non-decreasing in ζ, we can use bisection to search for the optimal ζ. By this, an optimal solution to (15) is obtained for any given c. Finally, according to Sections 4.2.5 and 11.4 of [30], the optimal c is solved by an outer-layer bisection search. Therefore, an optimal solution to Problem (13) is given by Algorithm 1. This completes the proof.

Appendix E
Proof of Lemma 2

Let us rewrite the functional h_1 as

$$h_1(z) = \frac{\int_0^\infty [y + z(y)]^2 d\mu_Y(y)}{\int_0^\infty [y + z(y)] d\mu_Y(y)}.$$

We need to prove that the functional h_1 is convex when restricted to any line that intersects its domain. For any $w \in L^2(\mu_Y)$, consider the function $u : \mathbb{R} \to \mathbb{R}$ defined as

$$u(\epsilon) = \frac{\int_0^\infty [z(y) + \epsilon w(y) + y]^2 d\mu_Y(y)}{\int_0^\infty [z(y) + \epsilon w(y) + y] d\mu_Y(y))}$$

with domain

$$\textbf{dom}\ u = \{\epsilon : z(y) + \epsilon w(y) \in [0, M],\ \forall y \geq 0, \epsilon \in \mathbb{R}\}.$$

Since the function $\epsilon \to [z(y) + \epsilon w(y) + y]^2$ is convex, the function $x \to \{[z(y) + (\epsilon + x)w(y) + y]^2 - [z(y) + \epsilon w(y) + y]^2\}/x$ is non-decreasing and bounded from above on $(0, 1]$. By using monotone convergence theorem [33, Theorem 1.5.6], we can interchange the limit and integral operators such that

$$\frac{d}{d\epsilon} \int_0^\infty [z(y) + \epsilon w(y) + y]^2 d\mu_Y(y) = \int_0^\infty 2[z(y) + \epsilon w(y) + y] w(y) d\mu_Y(y).$$

Similarly,

$$\frac{d}{d\epsilon} \int_0^\infty [z(y) + \epsilon w(y) + y] d\mu_Y(y) = \int_0^\infty w(y) d\mu_Y(y).$$

By this, we have

$$\frac{du}{d\epsilon} = \frac{\int_0^\infty 2[z(y) + \epsilon w(y) + y] w(y) d\mu_Y(y)}{\int_0^\infty [z(y) + \epsilon w(y) + y] d\mu_Y(y)} - \frac{\int_0^\infty [z(y) + \epsilon w(y) + y]^2 d\mu_Y(y) \int_0^\infty w(y) d\mu_Y(y)}{\left[\int_0^\infty [z(y) + \epsilon w(y) + y] d\mu_Y(y)\right]^2}. \tag{60}$$

After some additional manipulations, we can obtain

$$\frac{d^2u}{d^2\epsilon} = \frac{2\left[\int_0^\infty [y + z(y)] d\mu_Y(y) \int_0^\infty w(y) d\mu_Y(y)\right]^2}{\left[\int_0^\infty [z(y) + \epsilon w(y) + y] d\mu_Y(y)\right]^3} \times \int_0^\infty \left[\frac{y + z(y)}{\int_0^\infty [y + z(y)] d\mu_Y(y)} - \frac{w(y)}{\int_0^\infty w(y) d\mu_Y(y)}\right]^2 d\mu_Y(y).$$

Since $z(y) + \epsilon w(y) \geq 0$ for all y on **dom** u, we have $\frac{d^2u}{d^2\epsilon} \geq 0$. Hence, the function u is convex for all $w \in L^2(\mu_Y)$. By this, the functional h_1 is convex, which completes the proof.

Appendix F
Proof of Theorem 4

The Lagrangian of Problem (19) is determined as

$$L_1(z, \zeta, \gamma, \tau) = \frac{\int_0^\infty [y + z(y)]^2 d\mu_Y(y)}{2\int_0^\infty [y + z(y)] d\mu_Y(y)} + \zeta\left[1/f_{\max} - \int_0^\infty [y + z(y)] d\mu_Y(y)\right] - \int_0^\infty \gamma(y) z(y) d\mu_Y(y) + \int_0^\infty \tau(y)(z(y) - M) d\mu_Y(y),$$

where $\zeta \in \mathbb{R}$, $\gamma, \tau \in L^2(\mu_Y)$ are dual variables. According to [31, Proposition 3.3.2] and [32, pp. 70-72], the KKT theorem remains valid for the Lebesgue space $L^2(\mu_Y)$. Hence, if a vector (z, ζ, γ, τ) satisfies the KKT conditions (61)-(67), it is an optimal solution to (19). The KKT conditions are given by:

$$z = \min_{x \in L^2(\mu_Y)} L_1(x, \zeta, \gamma, \tau), \tag{61}$$

$$\zeta \geq 0, \int_0^\infty [y + z(y)] d\mu_Y(y) \geq 1/f_{\max}, \tag{62}$$

$$\gamma(y) \geq 0, z(y) \geq 0, \forall\ y \geq 0, \tag{63}$$

$$\tau(y) \geq 0, z(y) \leq M, \forall\ y \geq 0, \tag{64}$$

$$\zeta \left[1/f_{\max} - \int_0^\infty [y + z(y)] d\mu_Y(y)\right] = 0, \tag{65}$$

$$\gamma(y) z(y) = 0, \forall\ y \geq 0, \tag{66}$$

$$\tau(y)(z(y) - M) = 0, \forall\ y \geq 0. \tag{67}$$

We now solve the KKT conditions by using the calculus of variations. For any fixed (ζ, γ, τ), the Gâteaux derivative of the Lagrange L_1 in the direction of $w \in L^2(\mu_Y)$ at $z \in L^2(\mu_Y)$ is defined as

$$\delta L_1(z; w, \zeta, \gamma, \tau) \triangleq \lim_{\epsilon \to 0} \frac{L_1(z + \epsilon w, \zeta, \gamma, \tau) - L_1(z, \zeta, \gamma, \tau)}{\epsilon}.$$

Similar to the derivations of (60), we can obtain

$$\begin{aligned}
&\delta L_1(z; w, \zeta, \gamma, \tau) \\
&= \int_0^\infty \Bigg[\frac{y + z(y)}{\int_0^\infty [y + z(y)] d\mu_Y(y)} - \frac{\int_0^\infty [y + z(y)]^2 d\mu_Y(y)}{2\left[\int_0^\infty [y + z(y)] d\mu_Y(y)\right]^2} \\
&\quad - \zeta - \gamma(y) + \tau(y) \Bigg] w(y) d\mu_Y(y), \quad \forall\ w \in L^2(\mu_Y).
\end{aligned}$$

Then, $z(\cdot)$ is an optimal solution to (61) if and only if [34, p. 710]

$$\delta L_1(z; w, \zeta, \gamma, \tau) \geq 0, \quad \forall\ w \in L^2(\mu_Y).$$

By $\delta L_1(z; w, \zeta, \gamma, \tau) = -\delta L_1(z; -w, \zeta, \gamma, \tau)$, we deduce

$$\delta L_1(z; w, \zeta, \gamma, \tau) = 0, \quad \forall\ w \in L^2(\mu_Y).$$

Since $w(\cdot)$ is arbitrary, we have

$$\begin{aligned}
&\frac{y + z(y)}{\int_0^\infty [y + z(y)] d\mu_Y(y)} - \frac{\int_0^\infty [y + z(y)]^2 d\mu_Y(y)}{2\left[\int_0^\infty [y + z(y)] d\mu_Y(y)\right]^2} \\
&- \zeta - \gamma(y) + \tau(y) = 0, \ \forall\ y \geq 0.
\end{aligned} \tag{68}$$

For notational simplicity, let us define

$$\beta \triangleq \zeta \int_0^\infty [y + z(y)] d\mu_Y(y) + \frac{\int_0^\infty [y + z(y)]^2 d\mu_Y(y)}{2 \int_0^\infty [y + z(y)] d\mu_Y(y)}. \tag{69}$$

Since $\mathbb{E}[Y] > 0$, we have $\beta > 0$. The optimal primal solution $z(\cdot)$ is obtained by considering the following three cases:

Case 1: If $\gamma(y) = \tau(y) = 0$, then by (68) and (69), we obtain $z(y) = \beta - y$. In this case, we require $\beta - y \in [0, M]$ by (63) and (64).

Case 2: If $\gamma(y) > 0$, then by (66), $z(y) = 0$.

Case 3: If $\tau(y) > 0$, then by (67), $z(y) = M$.

In summary, the optimal primal solution $z(\cdot)$ is given by (21).

The optimal dual variable β is obtained by considering two cases:

Case 1: $\zeta > 0$. Then, (65) and (69) imply that

$$\int_0^\infty [y + z(y)] d\mu_Y(y) = \frac{1}{f_{\max}}, \beta \geq \frac{\int_0^\infty [y + z(y)]^2 d\mu_Y(y)}{2 \int_0^\infty [y + z(y)] d\mu_Y(y)}. \tag{70}$$

Case 2: $\zeta = 0$. Then, (62) and (69) imply that

$$\int_0^\infty [y + z(y)] d\mu_Y(y) \geq \frac{1}{f_{\max}}, \beta = \frac{\int_0^\infty [y + z(y)]^2 d\mu_Y(y)}{2 \int_0^\infty [y + z(y)] d\mu_Y(y)}. \tag{71}$$

Combining (70) and (71), yields

$$\int_0^\infty [y + z(y)] d\mu_Y(y) = \max\left(\frac{1}{f_{\max}}, \frac{\int_0^\infty [z(y) + y]^2 d\mu_Y(y)}{2\beta}\right).$$

Then, (22) is obtained. This completes the proof.

Appendix G
Proof of Theorem 5

Since $\mathbb{E}[Y] \geq 1/f_{\max}$, the constraint (20) is always satisfied and can be removed. If the zero-wait policy is optimal, we solve the KKT conditions of Problem (19) without the constraint (20). By this, we can obtain that the optimal primal solution is given by (21) almost everywhere, and the optimal dual solution β must satisfy

$$\beta \leq y_{\inf}, \quad \beta = \frac{\mathbb{E}[Y^2]}{2\mathbb{E}[Y]}, \tag{72}$$

from which (25) follows.

Next, we prove the reverse direction. If (25) holds, by $\mathbb{E}[Y] > 0$, we can get $\mathbb{E}[Y^2] > 0$ and $y_{\inf} > 0$. Hence, there must exist a $\beta > 0$ satisfying (72). Substituting (72) and $\mathbb{E}[Y] \geq 1/f_{\max}$ into Theorem 4, yields that the zero-wait policy is optimal. This completes the proof.

Appendix H
Proof of Lemma 3

1). When the correlation coefficient between Y_i and Y_{i+1} is -1, $Y + Y'$ is equal a constant value with probability one. Choosing $z(y) = 0$, $\zeta = 0$, $c = g(Y + Y')$, $\gamma(y) = \tau(y) = 0$, one can show that the KKT conditions (43)-(49) are satisfied.

2). If the Y_i's are equal to a constant value, $Y + Y'$ is equal a constant value. The remaining proof follows from part 1).

3). When $g(\cdot)$ is a constant function, all policies are optimal. This completes the proof.

References

[1] A. R. Al-Ali, I. Zualkernan, and F. Aloul, "A mobile GPRS-sensors array for air pollution monitoring," *IEEE Sensors Journal*, vol. 10, no. 10, pp. 1666–1671, Oct 2010.

[2] T. Gandhi and M. M. Trivedi, "Pedestrian protection systems: Issues, survey, and challenges," *IEEE Transactions on Intelligent Transportation Systems*, vol. 8, no. 3, pp. 413–430, Sept 2007.

[3] V. Terzija, G. Valverde, D. Cai, P. Regulski, V. Madani, J. Fitch, S. Skok, M. M. Begovic, and A. Phadke, "Wide-area monitoring, protection, and control of future electric power networks," *Proceedings of the IEEE*, vol. 99, no. 1, pp. 80–93, Jan 2011.

[4] S. Kaul, R. Yates, and M. Gruteser, "Real-time status: How often should one update?" in *Proc. IEEE INFOCOM*, Apr. 2012, pp. 2731–2735.

[5] ——, "Status updates through queues," in *Conf. on Info. Sciences and Systems*, Mar. 2012.

[6] R. Yates and S. Kaul, "Real-time status updating: Multiple sources," in *Proc. IEEE Int. Symp. Inform. Theory*, Jul. 2012.

[7] C. Kam, S. Kompella, and A. Ephremides, "Age of information under random updates," in *Proc. IEEE Int. Symp. Inform. Theory*, 2013.

[8] ——, "Effect of message transmission diversity on status age," in *Proc. IEEE Int. Symp. Inform. Theory*, June 2014, pp. 2411–2415.

[9] M. Costa, M. Codreanu, and A. Ephremides, "Age of information with packet management," in *Proc. IEEE Int. Symp. Inform. Theory*, June 2014, pp. 1583–1587.

[10] L. Huang and E. Modiano, "Optimizing age-of-information in a multi-class queueing system," in *Proc. IEEE Int. Symp. Inform. Theory*, 2015.

[11] R. Yates, "Lazy is timely: Status updates by an energy harvesting source," in *Proc. IEEE Int. Symp. Inform. Theory*, 2015.

[12] T. Bacinoglu, E. T. Ceran, and E. Uysal-Biyikoglu, "Age of information under energy replenishment constraints," in *Proc. Info. Theory and Appl. Workshop*, Feb. 2015.

[13] S. M. Ross, *Applies Probability Models with Optimization Applications*. San Francisco, CA: Holden-Day, 1970.

[14] H. Mine and S. Osaki, *Markovian Decision Processes*. New York: Elsevier, 1970.

[15] D. Hayman and M. Sobel, *Stochastic models in Operations Research, Volume II: Stochastic Optimizations*. New York: McGraw-Hill, 1984.

[16] E. A. Feinberg, "Constrained semi-markov decision processes with average rewards," *Zeitschrift für Operations Research*, vol. 39, no. 3, pp. 257–288, 1994.

[17] D. P. Bertsekas, *Dynamic Programming and Optimal Control*, 3rd ed. Belmont, MA: Athena Scientific, 2005, vol. 1.

[18] R. Bellman, *Dynamic Programming*. Princeton University Press, 1957.

[19] M. J. Neely, "Dynamic optimization and learning for renewal systems," *IEEE Trans. Auto. Control*, vol. 58, no. 1, pp. 32–46, Jan 2013.

[20] A. Federgruen, P. J. Schweitzer, and H. C. Tijms, "Denumerable undiscounted semi-Markov decision processes with unbounded rewards," *Mathematics of Operations Research*, vol. 8, no. 2, pp. 298–313, 1983.

[21] D. Klabjan and D. Adelman, "Existence of optimal policies for semi-Markov decision processes using duality for infinite linear programming," *SIAM J. Control Optim.*, vol. 44, no. 6, pp. 2104–2122, Jan. 2006.

[22] F. J. Beutler and K. W. Ross, "Optimal policies for controlled Markov chains with a constraint," *Journal of Mathematical Analysis and Applications*, vol. 112, no. 1, pp. 236 – 252, 1985.

[23] ——, "Time-average optimal constrained semi-Markov decision processes," *Adv. in Appl. Probab.*, vol. 18, no. 2, pp. 341–359, 1986.

[24] M. Baykal-Gürsoy and K. Gürsoy, "Semi-Markov decision processes: nonstandard criteria," *Probability in the Engineering and Informational Sciences*, vol. 21, pp. 635–657, Oct. 2007.

[25] E. A. Feinberg, "Constrained discounted semi-Markov decision processes," in *Markov Processes and Controlled Markov Chains*, Z. Hou, J. A. Filar, and A. Chen, Eds. Kluwer Academic Publishers, 2002.

[26] E. Altman, *Constrained Markov decision processes*. Boca Raton, FL, USA: Chapman & Hall/CRC, 1999.

[27] C. Derman, *Finite State Markovian Decision Processes*. Orlando, FL, USA: Academic Press, Inc., 1970.

[28] M. J. Neely, *Stochastic Network Optimization with Application to Communication and Queueing Systems*. Morgan & Claypool, 2010.

[29] W. Rudin, *Real and Complex Analysis*, 3rd ed. New York: McGraw-Hill Book Co., 1987.

[30] S. Boyd and L. Vandenberghe, *Convex Optimization*. Cambridge, U.K.: Cambridge Univerisity Press, 2004.

[31] D. Butnariu and A. N. Iusem, *Totally Convex Functions for Fixed Points Computation and Infinite Dimensional Optimization*. Norwell, MA: Kluwer Academic Publisher, 2000.

[32] G. G. Magaril-ll'yaev and V. M. Tikhomirov, *Convex Analysis: Theory and Applications*. American Mathematical Society, 2003.

[33] R. Durrett, *Probability: Theory and Examples*, 4th ed. Cambridge Univerisity Press, 2010.

[34] D. P. Bertsekas, *Nonlinear Programming*, 2nd ed. Belmont, MA: Athena Scientific, 1999.